

REPORT DOCUMENTATION PAGE

*Form Approved
OMB No. 0704-0188*

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE	10-22-2003	2. REPORT TYPE	FINAL	3. DATES COVERED (From - To)	05/03/00 – 12/31/02
4. TITLE AND SUBTITLE			A Software Control Framework for Learning Coordinated, Multi-Robot Strategies in Open Environments		
			5a. CONTRACT NUMBER DABT63-00-1-0004		
			5b. GRANT NUMBER DABT63-00-1-0004		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Roderic Grupen			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)			8. PERFORMING ORGANIZATION REPORT NUMBER 01		
Computer Science Department University of Massachusetts 140 Governors Drive Amherst, MA 01003			10. SPONSOR/MONITOR'S ACRONYM(S) DARPA		
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
Defense Advanced Research Projects Agency 3701 North Fairfax Drive Arlington, VA 22203-1714					

12. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release. Distribution is unlimited.

20031117 052

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The UMass MARS team contributed technology and infrastructure for the control of adaptive sensory and motor processes. We developed a suite of techniques for capturing interesting process dynamics in specific run-time contexts in order to learn control decisions. Particular attention was paid to the ability of teams of robots to adapt dynamically to changes in environment and mission requirements. The UMass effort marries high-level process descriptions, discrete event analysis and model checking, learning and stochastic exploration, and a control theoretic substrate to accomplish these goals.

Distributed control Technologies have been transferred to SPAWAR in San Diego, CA. "Whole-body" distributed manipulation controllers and finger gaiting code for autonomous manipulation tasks have been ported to NASA-JSC for use in the Robonaut program. With related funding under the DARPA DASADA program, contractors of the U.S. Army for the Rotorcraft Pilot's Associate (RPA) and Theater High Altitude Area Defense (THAAD) programs are exploring the use of Containment Units (CUs) to build adaptable systems.

15. SUBJECT TERMS

process control, learning, distributed systems, grasping and manipulation

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Roderic Grupen
a. REPORT Unclassified	b. ABSTRACT Unclassified	c. THIS PAGE Unclassified	UL	11	19b. TELEPHONE NUMBER (include area code) (413)545-3280

Objective: The UMass MARS team contributed technology and infrastructure for the control of adaptive sensory and motor processes. In general, process dynamics continually change so we are developing systems that match run-time data to a family of process models - each of which captures interesting process dynamics in specific run-time contexts. These models accumulate over time and provide state estimates for learning algorithms that allocate sensory and motor resources and control. We also explore mechanisms for guaranteeing critical functional specifications during learning and in the resulting policy. Our goal was to develop a multi-level software architecture that allows a contingent of mobile platforms with varying sensory apparatus, mobility, and computational power to cooperate while completing a multi-objective mission in an "open" environmental context. Particular attention was paid to the ability of teams of robots to adapt dynamically to changes in environment and mission requirements. A hybrid solution was sought that expresses a family of equivalent process control configurations that guarantee functional specifications and include event/feature monitors and contingencies for these events. The UMass effort marries high-level process descriptions, discrete event analysis and model checking, learning and stochastic exploration, and a control theoretic substrate to accomplish these goals.

Approach: We adopted a hierarchical approach in which plans and goals are represented and manipulated by high-level components and control is propagated down to the lower-level controllers. At the same time, run-time feedback from the reactive lower-level controllers is propagated upward to establish the state in the high-level deliberative processes. We assume that such knowledge is not, in general, available *a priori*. In this manner, structure is imposed from above on the range of reactive responses permitted and models of the run-time environment are delivered from below. Reconfiguration decisions are produced at the level of the hierarchy where adequate control knowledge and resources exist to respond to inadequate performance.

The bi-directional report/react responsibilities above are expressed in the form of a software construct called the "Containment Unit" (CU). CUs "contain" faults, i.e. they preserve a specific set of properties during the execution of a sensorimotor task - like safety, precision, etc. They do so by using whatever resources have been allocated to them and report both results and context (observable, but uncontrollable conditions) to their clients. They insulate higher-level processes from the details of the run-time configuration and afford lower-level processes a bounded degree of flexibility in satisfying mission objectives. The CU is also the place where policies are acquired to manage resources in order to optimize the behavior of the run-time system with respect to the CU specification.

The selection and activation of containment units is performed by a process modeling and planning component that can assign goals and resources based on the overall mission objectives. It can do so using the abstraction provided by the CU. The CU can be viewed as a subtask - expressing a kind of "divide-and-conquer" approach with which to constrain the complexity associated with learning resource allocation policies.

The UMass effort consists of several key components.

Multi-Agent Systems - provides high-level guidance and deliberation where appropriate in place of stochastic search.

Our effort is focused on applying the MAS scheduling work to a 3 sensor node configuration of the ANTS millimeter-length radar platform. Progress towards a new agent level negotiation protocol for sensor head allocation is noteable. In addition to hardware tests, this new protocol was evaluated empirically in two simulation environments. Our agent architecture is now more robust when operating in a real-world environment. Our organizational structure has been augmented to support task migration. Specifically, the goal responsible for tracking a target can now migrate between agents to better maintain spatial locality. Addressing robustness, a mechanism was added allowing agents to easily propagate liveness information. This was done by using historical directory service query information to direct the flow of new or modified data. We have simulated a mechanism allowing agents to form the clusters dynamically to segment information flow in the system. We are in the process of exploring how this technique can be adapted to our implementation.

Most recently, we have created an eight-node hardware layout, and are currently evaluating our architecture in this context.

Publications:

- R. Vincent, B. Horling, V. Lesser, and T. Wagner. Implementing Soft Real-Time Agent Control. In Proceedings of the 5th International Conference on Autonomous Agents, pp. 355-362, ACM Press, Montreal, June 2001.
- B Horling, R.Vincent, R. Mailler, J. Shen, R. Becker, K. Rawlins, and V. Lesser. Distributed Sensor Network for Real Time Tracking. In Proceedings of the 5th International Conference on Autonomous Agents, pp. 417-424, Montreal, June 2001, ACM Press.
- Raja, Anita, Lesser, Victor. Meta-Level Control in Multi-Agent Systems. UMass Computer Science Technical Report 01-49, November, 2001.
- Mailler, Roger, Vincent, Regis, Lesser, Victor, Middlekoop, Tim, and Shen, Jiaying. Soft-Real Time, Cooperative Negotiation for Distributed Resource Allocation. In AAAI for Symposium on Negotiation, Falmouth, MA, November, 2001.

Process Description - abstract descriptions of the “process” of reconfiguration, including the instantiation of subordinate processes, the logging of event streams, the estimation of state for process control, the reconfiguration of subordinate processes, and the generation of an event stream to superordinate processes.

Our work on Little-JIL is focused on maturing the concept and implementation of containment units. An abstract specification of a CU has been defined, and Little-JIL has been used to build and analyze, and execute several CUs including the TRACK

HUMAN CU mentioned in previous reports, and an ILLUMINATION CONTROL CU which uses COTS sensors and actuators to control and monitor lamps and respond to failures.

The TRACK HUMAN containment unit employs Reinforcement Learning (RL) to learn a policy for selecting regions of interest as an alternative to exhaustive matching, and using Timed Automata as a vehicle for specifying and verifying properties of low-level controllers. We have constructed a CU that follows a subject using motion sensors and predicts which sensors will be detect the subject next. This CU uses the resource collection features being added to the language and will allow us to experiment with resource allocation algorithms in containment units. This CU also uses the reaction mechanism in Little-JIL which is new development and work has begun on defining the semantics of this feature and specifying how it interacts with the proactive language elements.

In addition to the work on resource collections, the resource manager has received a significant amount of work – our initial prototype has been replaced with a new implementation using Infonyte-DB from Infonyte GmbH. This implementation gives us access to the powerful XQL query language for specifying resource requirements in Little-JIL process programs. Resource management is a central part of the containment unit architecture.

Publications:

- L. Osterweil, A. Wise, J. Cobleigh, L. Clarke and B. Lerner. Architecting Dynamic Systems Using Containment Units. In Working Conference on Complex and Dynamic Systems Architecture, Brisbane, Australia, December 2001.
- B. Lerner, J. Cobleigh, L. Osterweil, A. Wise. Using Containment Units for Self Adaptation of Software. International Conference on Software Engineering, Buenos Aires, Argentina May 2002.

Sensory Processing - the creation and interpretation of signals to produce events or features of the controlled process at several levels of abstraction from the combination of control processes, sensors, and sensory abstraction.

An experimental platform has been assembled including an array of sensor type and mobile platforms that solve sensory and motor tasks by forming virtual robots from networked services. Computation is likewise distributed via radio ethernet and analog RF. This design integrates perception, control, computation and communication into a hierarchical organization that supports distributed autonomous behavior. Details are provided at the UMass MARS website. The sensor array platform has been developed to support multi-subject tracking using a single panoramic camera, and work has continued on using a camera mounted on a mobile robot.

For multi-subject tracking, we developed a method that acquires new features as needed to extract hidden state and support control decisions. One such approach annotates subjects with color-histograms to track multiple subjects when trajectories cross and

to recognize the same subject when subjects move across disjoint sensor populations. The same techniques support recovery from momentary loss of track due to occlusions or detection failures.

Publications:

- R. Uppala, D. Karuppiah, M. Brewer, S. Ravela , and R. Grupen. "On Viewpoint Control." International Conference on Robotics and Automation, Washington DC, May 2002.
- G. Holness, D. Karuppiah, S. Uppala, S. Ravela, and R. Grupen. "A Service Paradigm for Reconfigurable Agents." In Second International Workshop on Infrastructure for Agents, MAS, and Scalable MAS, Montreal, May 2001.
- S. Ravela and A. Hanson. "On Multi-Scale Differential Features for Face Recognition." In Vision Interface 01, Ottawa, June 2001.
- D. Karuppiah, Z. Zhu, P. Shenoy, and E. Riseman. "A Fault-Tolerant Distributed Vision System Architecture for Object Tracking in a Smart Room." In International Workshop on Computer Vision Systems, Vancouver, July 2001.

Robotics - correct reactive control configurations, constraint maintenance in the context of functional specifications, control theoretic substrate, the expression of device and feedback control-level control knowledge, mechanisms for learning and deploying correct families of control alternatives that limit the scope of stochastic exploration and policy formation, the construction and reconfiguration of virtual sensory and motor systems.

We have experimented with a family of so-called pick-and-place problems that introduce challenging constraints to motion planning and are not adequately addressed using current technologies. We have produced results with a learning system on our humanoid platform that learns re-useable skills in a developmental sequence and that augments its state representation to model tasks. Our robot learns (via reinforcement learning) to anticipate constraints that only become apparent in the near future in order to make immediate grasping decisions. For example, in order to fit an object into a hole, the robot must learn to pick up the object in such a way that the grip itself does not interfere with the insertion operation. Finally, we have demonstrated a learning technique which enables the robot to make accurate head orienting movements in response to audio stimuli (as sensed by a quadraphonic audio system). We take advantage of the fact that audio and visual stimuli are often co-occurring. A visual-based motion tracking controller not only enables the robot to follow moving stimuli in its environment, but we show how the requisite movements can also be used to calibrate movements that are made in response to audio stimuli that are not in the camera's field of view.

Publications:

- Coelho, J. , Piater, J. and Grupen, R., "Developing Haptic and Visual Perceptual Categories for Reaching and Grasping with a Humanoid Robot, Journal of

Robotics and Autonomous Systems," special issue on Humanoids, Volume 37, 2001, pp. 195-218.

- Grupen, R. and Coelho, J., "Acquiring State from Control Dynamics to Learn Grasping Policies for Robot Hands," Robotics Society of Japan (RSJ) International Journal on Advanced Robotics.
- Huber, M., Grupen, R., "Robust Finger Gaits from Closed-Loop Controllers," ICRA02, October, 2001.
- Wheeler, Fagg, Grupen, "Learning Prospective Manipulation Strategies," International Conference on Learning and Development, June, 2002.

Human Interface - Wearables and Augmented/Virtual Reality Our work now encompasses interface and supervision of multi-robot teams by a human operator or client. We have developed a method for importing the running state of a mobile robot that will be used to render a 3D model of the robots' environment, which in turn allows for a remote user to explore the environment from a safe distance. This is a first step towards allowing the user to make high-level requests of the robot array (e.g., to task a set of robots to explore specific areas or provide updates to the current model) from within the virtual environment. We are also exploring the use of user intervention as a means of providing training experience to the robot control system.

We have also developed a network-distributed object representation system. This system enables the scattering of three-dimensional objects across a set of servers located on a common network. Objects may directly represent the state of sensors or a summary of data that has been obtained from a set of sensors (for example, the three-dimensional location of a tracked subject). Because these objects are network accessible in a uniform manner, the construction of new virtual environments from data derived from different robots and sensor state is efficient and simple.

Publications:

- Amstutz, P. and Fagg, A. H. (2002), Real Time Visualization of Robot State with Mobile Virtual Reality, International Conference on Robotics and Automation (ICRA'02)
- Davis, J. A., Fagg, A. H., Levine, B. N. (2001), Wearable Computers as Packet Transport Mechanisms, In Highly-Partitioned Ad-Hoc Networks, Proceedings of the International Symposium on Wearable Computing, Zurich, Switzerland, October 2001.

Accomplishments: During fiscal years 2001 and 2002, the effort has produced significant deliverable technologies.

1. Network time protocols for synchronization across a network of distributed sensor and motor services. We developed a mechanism to save synchronized images sequences from two cameras to create reusable datasets and test our algorithms offline.

2. A common set of interfaces and object types for describing and controlling sensors and motors using Jini services and the JNI interface to embedded devices.
3. Demonstrations of virtual agents for guaranteeing bounded overwatch behavior wherein a designated subject (human and robotic) can be localized continuously in the presence of noise, occlusion, and hardware faults.
4. adaptive optimal, single robot path controllers have been verified in simulation for optimal energy and time performance in the context of generic path tracking tasks for a single robot,
5. demonstrated control expressions that compile into distributed implementations that satisfy global specifications and demonstrated how equivalent functional configurations vary with respect to time, energy, and communication bandwidth,
6. Developed integrated control algorithms for coordinated viewpoint control in mobile sensory systems. Each robot treats its partners as if they are stationary sensors and calculates a target location for themselves that minimizes uncertainty.
7. Created a formal description and library of histogram and texture feature vectors that aid in matching representations of objects derived from disparate sensors and feature vectors. New tools will shortly be available for determining if two disparate sensors refer to the same feature in the world even though their representations differ.
8. New algorithms for single subject tracking using an instantaneous triangulation uncertainty measure and multiple subject discrimination using color histograms.
9. An assortment of virtual sensors combining panoramic, acoustic, pyroelectric, millimeter-wave radar, and pan-tilt-zoom visual systems in a combination of stationary and mobile configurations.
10. The Little-JIL process programming language for guiding and analyzing the behavior of autonomous robotic resources. Our Little-JIL execution environment has been used to execute the "Illumination Control" CU, and as noted in the 2001 Annual Report, demonstrated at the DASASA 2001 PI Meeting in June.
11. Prototypes of human interfaces including wearable/immersive forms to support human-robot collaboration,
12. Multi-robot, coordinated search algorithms have been described completely, verified in simulation, implemented on our uBot platforms and ported to SPAWAR, San Diego.
13. Several containment units have been designed and documented for creating virtual agents to guarantee bounded overwatch behavior wherein a designated subject (human and robotic) can be localized continuously in the presence of noise, occlusion, and hardware faults. This has been demonstrated for tracking a single moving object from a camera pair made up of a mobile and a fixed camera.

14. Multi-objective path planning to plan motion for mobile sensor platforms that can simultaneously maximize visibility of a tracked subject, minimize the triangulation errors, and manage obstacle avoidance.
 - (a) search and mapping behavior,
 - (b) leader-follower behavior,
 - (c) a multi-robot behavior for preserving network connectivity among peers in a co-ordinated task, and
 - (d) a bounded-overwatch behavior that localizes individuals within a coordinated team using virtual stereo.
15. Implemented prototype human interfaces including wearable/immersive forms to support human-robot collaboration Based upon these technologies, we have implemented the following demonstrations:
 - (a) visualization of robot state: the kinematic state of the UMass Torso is depicted in real time in a virtual environment;
 - (b) “X-ray vision” the user is able to perceive moving objects through occluding walls and other objects; and
 - (c) a dynamic display of distributed robot search task with real-time updates of robot positions, target objects, and obstacles.
16. We have completed a beta version of a communication system that supports the distributed representation of three-dimensional environments, whose state can be updated using sensor streams and can be shared among multiple clients.
17. Developed and standardized JINI Sensor and Motor Services. Sensors and motors register their services in a lookup and CUs are configured by downloading the service proxy from the lookup that subscribes to the required event streams.
18. Implemented “whole-body” distributed grasping behavior and ported finger gaiting code for autonomous manipulation tasks to NASA-JSC.
19. Developed a “developmental programming paradigm” that compiles control knowledge incrementally in a developmental sequence. Generated quantitative performance impact of the approach in distributed control tasks. Yielded double the performance in half the learning time for quasistatic walking tasks and visual discrimination tasks.
20. Implemented a new approach to recovering hidden state information with interaction-based state estimation - an extension to traditional HMM-based approaches.
21. Constructed a computational account of developmental dynamics in infants learning to solve manipulation problems combining vision and touch. Experimental results on robots reveal the same dynamics as observed in human infants.

The following text discusses our progress in more detail and on several dimensions of the project.

Hardware Infrastructure - An experimental platform has been assembled including an array of sensor type and mobile platforms that solve sensory and motor tasks by forming virtual robots from networked services. Computation is likewise distributed via radio ethernet and analog RF. This design integrates perception, control, computation and communication into a hierarchical organization that supports distributed autonomous behavior. Details are provided at the UMass MARS website.

Coordinated Motion Controllers - In collaboration with the UMass SDR project, 4 motion platforms (UMASS UBots) have been prototyped consisting of vision (and future acoustic) sensors, IR obstacle detection, future capacity for IRDA broadcast communication channel, a DIMMPC controller with RT kernel, and I2C communications between embedded Scenix subsystems. An RF rangeLAN point-to-point communication facility for robots to communicate with peers has been implemented.

Coordinated policies for search and mapping tasks have been demonstrated using containment units for localizing participants using a bounded overwatch and for preserving network connectivity. These CUs include parametric "push" and "pull" constructs to preserve critical kinematic relationships.

Viewpoint Planning - A CU for tracking moving generic features over time has been constructed that supports a variety of localization and identification tasks. Correspondences over multiple sensor modalities can often be more discriminating than single modalities. For example, a human subject presents a conjunction of thermal, visual, motion, and acoustic signatures. Our UBot is yet another distinctive combination of these features so that policies can be selected for tracking that also address identification goals.

Moreover, the precision of localization depends on the resources employed, the virtual stereo geometry, network latency, etc. The instantaneous uncertainty of networked observers can be minimized by evaluating the kinematics of the virtual sensor. We have demonstrated sequences of observers that minimize observation uncertainty and that provide a seamless stream of spatial reports even as the subject moves through occlusions and in the presence of hardware faults.

JINI Sensor and Motor Services - To organize the services required to implement a particular virtual robot configuration, the Jini service protocol has been adopted. Sensors and motors register their services in a lookup and CUs are configured by downloading the service proxy from the lookup that subscribes to the required event streams. The event stream can consist of features on signals, spatial data types constructed by teams of sensors, features derived from the dynamics of sensorimotor processes, or discrete events such as hardware faults. A network time protocol is used to synchronize services and the client CU interpolates between samples to improve performance. The resulting

spatial acuity of visual trackers constructed using these techniques has been quantified. We have achieved precision on the order of centimeters when tracking human subjects moving at moderate velocities in typical interior settings.

Little-JIL Process Language and CU compilation - Little-JIL is a graphical language for defining processes that reconfigure the resources engaged by an agent to achieve a specified task. Little-JIL provides a high-level programming language for the specification and analysis of Containment Units (CUs). Several demonstrations of this capability were shown at the DASADA 2001 PI meeting held June 4-7, 2001 in Baltimore MD.

Developmental Programming Paradigm - Examples of the programming methodology whereby simple skills precede peer-to-peer coordination strategies which, in turn, precede team-oriented collaborations have been constructed for distributed search tasks. The approach depends on compiling control knowledge incrementally in a developmental sequence. Quantifiable results of the impact of this approach on learning in distributed systems are expected during the next quarter.

Experiments with interaction-based state estimation on which to base generalizable control decisions have been demonstrated in simulation. Control state is estimated by observing the transient response of a discrete set of linear controllers that collectively identify the operational context. Thus far, results confirm the hypothesis that these representations generalize to novel run-time situations. A report detailing the results obtained for time, energy, and precision objectives in path controllers on our UMass UBot platforms is currently being prepared.

Current Plan: We plan to continue to develop these technologies with a heavy emphasis on integration and demonstration.

1. Networked Team Services - A large class of indoor search and inspection tasks can be implemented by coordinating the activity of groups of 5 or so platforms configured to cover local, multiple room regions in a floorplan. Hierarchical CUs for a coordinated team will be constructed that manage the high-level roles of the robots (search, comm, etc), to maintain critical pre-requisites of this plan (network topology and connectivity), and to manage the prerogatives of individuals in the team (power, time, etc). This three-level command hierarchy will consist at policies learned at three independent levels: those that define the roles each robot will occupy, those that maintain QOS in peer-to-peer interactions, and finally those that express the orthogonal interests of the individual participants.
2. Develop Additional Basic Services - acoustic localization and identification, extraction of signature features with which to disambiguate multiple subjects - color, shape, and appearance models, human id CUs, activity maps and other abstractions of human (and

robot) behavior that can be used to recognize human subjects, to facilitate resource (re)configuration, and to support multi-subject tracking.

3. Collaborative Human-Robot Interfaces - CUs that service and maintain user interfaces and CUs that employ human services in response to other clients are essential to effective collaborative interfaces. We are constructing a 3D graphics engine to augment a visual scene and/or to render virtual environments. The target computational and display device is a Xybernaut wearable computer that can render a visual interface on a heads up display at approximately 10 frames per second. The *track subject* CU is used to localize the human and orientation is determined by an accelerometer/gyro combination. In the coming year, we plan to realize interfaces that monitor and summarize the activity of distributed teams of robots and that permit the human to participate in sensorimotor tasks in a variety of ways.

Much of this work will be conducted in the context of our follow-up MARS contract with Johnson Space Center, Vanderbilt, USC, and MIT. The MARS contract expired officially after a no-cost extension in December of 2002.

Technology Transfer: One substantial technology transition result of this research project is education of students. Twelve students are involved in our MARS project. They are trained to contribute to technology transfer when they move to industrial or government labs upon graduation.

Our coordinated motion controller have been ported to simulators at SPAWAR in San Diego, CA.

Implemented "whole-body" distributed grasping behavior and ported finger gaiting code for autonomous manipulation tasks to NASA-JSC.

We are continuing to identify opportunities for the direct transition of technology under the project to DOD applications. We recently acquired a small NASA contract to deliver our controllers for applications to the Robonaut program.

With related funding under the DARPA DASADA program, contractors of the U.S. Army for the Rotorcraft Pilot's Associate (RPA) and Theater High Altitude Area Defense (THAAD) programs are exploring the use of Containment Units (CUs) to build adaptable systems.